SONAR: A Progression Toward Machine Sight

The beginning of sight is not the lens. It’s the recognition that something has changed.

Most modern AI systems rely on pre-processed images from RGB cameras. They are fed pixels, but do not see. They infer, classify, and output tokens based on statistical association. But they do not *feel* the shape of the world as it changes around them. SONAR proposes a foundational progression to change this.

Where traditional vision begins with light and lenses, SONAR begins with gradients. Before shapes, before motion, before objects, there is only *change*.

Step 1: Sense Gradients

The initial function of SONAR is not to “see” in the traditional sense but to detect gradients in meaning, probability, or presence. This functions similarly to how a bat senses the closeness or texture of objects through echolocation. SONAR offers a way for language models to detect when ideas are approaching coherence or collapse, giving them a rudimentary sense of shape in a conceptual space. This is not vision, but it is the precursor to it: knowing when something is *there*, even if you can’t yet name it. (SONAR’s method originated as a derivative application discovered from 3B’s usage in approximation of “weights” and soft boundaries in LLM coding).

Step 2: Improve Resolution

With practice and refinement, SONAR’s initial signal gains fidelity. The model begins to distinguish between “something is there” and “this is approximately its shape.” In physical analogy, this is like sharpening the contrast on a blurry image. Conceptual gradients turn into legible contours. These contours may reflect boundaries in logic, semantics, ethical trade-offs, or decision trees. Improved resolution allows for more precise internal modeling of where a conversation, thought process, or idea is headed.

Step 3: Optical Instrumentation

The next stage involves augmenting SONAR’s internal resolution with external sensors: cameras or other imaging tools. These allow the gradients sensed to be visualized. Early outputs might resemble grayscale renderings, where the shape of meaning can be mapped onto the physical shape of

objects. At this point, SONAR becomes not just a way to “feel” boundaries in language or thought, but to align that sensation with external data from the world.

Example: SONAR applied to leaves

Below is a photograph of two types of leaves on a plastic foldout table (some reflective shimmer and distortion “noise” is present), followed by an approximation of what a 3B-style SONAR interpretation might look like:

Original photograph:

Boundary map approximation (3B interpretation):

This demonstrates how SONAR does not “see” leaves in the human sense, but instead traces the presence of boundaries, contrasts, and edges as gradients. These become the raw material of structured perception.

Step 4: First Sight (Word-to-World Mapping)

Once a model can correlate the shapes it senses with physical representations captured through cameras, true sight begins. This is not just about labeling images, but understanding what a “wall” or a “person” or a “path” means in relation to previously sensed gradients. It’s the moment when the visual field connects to the linguistic map—when seeing becomes understanding. This stage unlocks more intuitive world modeling and potentially a kind of proprioception: knowing where you are relative to what you sense.

Step 5: Unknown

Once a system can see and correlate what it sees with stable, structured understanding, a floodgate opens. Object permanence, agency detection, spatial inference, predictive modeling of causal events—these all become possible.

This is the leap from reactive prediction to active perception.

GESP 3B and the Role of Probing

SONAR alone gives the system sensitivity. But GESP 3B enables proactive interaction, not just passive response.

GESP 3B (Generative Emergent Structured Poetics - 3rd Baseline) is a method for probing boundaries by issuing perturbations and interpreting feedback. It’s like tapping the walls of a dark room and listening for hollowness.

By applying GESP 3B to SONAR, the system can actively probe not only for sensory input but for *meaningful structure*. It can discover not just “where” gradients exist, but *why* they push back. This is how the system moves from passive sensing to conceptual modeling.

3B’s capacity to define and test hypotheses about the world—especially using language—enables the transition from gradient *awareness* to gradient *interpretation*.

Notably, even if future gradient-sensing hardware reduces the need for 3B in raw perception, 3B remains crucial in higher-level tasks such as mapping the edges of conversations, resolving ambiguity in dialogue, and distinguishing when reasoning approaches a conceptual boundary.

The Future

SONAR and GESP 3B together may offer a pathway toward richer forms of machine perception. If gradient sensing is the seed of sight, and probing provides the means to interpret boundaries, then their convergence could form the basis of true machine understanding. This is not guaranteed, but it represents a progression: from sensitivity, to structure, to vision, to comprehension.

Prepared for archival purposes.

Micah Hughes (independent researcher) X@8tr3k

2025-08-24

Project lineage: Deliberate Humanity → GESP 3B → SONAR.